Population Based Heuristics for the Aircraft Landing Problem

Hélène L. Pinol* John E. Beasley*

*The Business School, Imperial College London
South Kensington campus
London SW7 2AZ, United Kingdom
{h.pinol,j.beasley}@imperial.ac.uk

1 Introduction

Before actually landing at an airport, an aircraft must go through an approach stage directed by air traffic controllers. When entering the airport radar range, the aircraft's flight number, altitude and speed are transmitted to controllers within the air traffic control tower. Based on these information, controllers first direct the aircraft's pilot to an appropriate approach corridor. When the aircraft is closer to the airport, controllers give instructions for the speed and altitude of the aircraft in order to align it with the allocated runway. During peak hours, controllers must handle safely and effectively landings of a continuous flow of aircraft entering the radar range onto the assigned runway(s). The capacity of runways is highly constrained and this makes the scheduling of landings a difficult task to perform effectively. Because of environmental and geographic constraints, capacity cannot be simply increased by building new airport structures or runways. Therefore there is a need to develop effective decision support tools that provide useful help to controllers.

The decision problem that air traffic controllers face repeatedly over time is: what aircraft should land next and when should it land? As some airports have several runways, it is possible that more than one runway is available for landings. In that case, air traffic controllers must face an additional decision: what runway should be allocated to the next incoming aircraft?

In the single runway situation, the landing sequence adopted for incoming aircraft is often decided in a First-Come First-Served (FCFS) manner. The first aircraft that enters the radar range must land first, the second aircraft that enters lands second and so on. When several runways are available for landings, the FCFS scheduling is also often used such that aircraft land on their allocated runway in the corresponding order they appear in the radar range. The scheduling of an appropriate landing time for each aircraft is constrained by its characteristics. When an aircraft enters the radar range, a target (preferred) landing time is defined as the time the aircraft could land if it flies straight to the closest runway at its cruise (most economical) speed. This target landing time is bounded with an earliest landing time and a latest landing time. The earliest landing time is determined as the time the aircraft could land if it flies straight to the closest runway at its fastest speed with no holding. The latest landing time is determined as the time the aircraft could land if it is held for its maximal allowable time before landing. The time that elapses between the aircraft's earliest landing time and latest landing time is called its time window. Hence, scheduled landing times decided by air traffic controllers must lie within aircraft time windows, and these define a first set of constraints: time window constraints.

The flow of incoming aircraft is not homogenous, it contains different aircraft types. For the major UK airports such as London Heathrow and London Gatwick, aircraft are classified into five types: Light, Small, Lower-Medium, Upper-Medium and Heavy. These classes are mainly defined based on aircraft weight. All aircraft are creating wake vortices at the rear of the aircraft when landing. These vortices have a chaotic evolution and can cause serious turbulence to a closely following aircraft even to the extent of a crash. In order to maintain an aircraft's aerodynamic stability, separation distances based on the preceding aircraft types must be respected during landing. For example, if an Heavy aircraft is followed by an Upper-Medium aircraft, a separation distance of five nautical miles is requested. Whereas if an Heavy aircraft is following any kind of aircraft, the separation distance needed is only three nautical miles. In our research, separation distances are handled by converting them into separation times using a fixed landing speed depending on the aircraft type. These separation times are the main limiting factors on runway usage and, these define a second set of constraints: separation time constraints.

In addition to the above constraints, the aircraft landing problem (ALP) also includes an objective to assess how effective is the suggested use of the constrained capacity runway. In air traffic management, deviation from target landing times is a key factor. It is also useful to have the ability to change the objective regarding air traffic density or weather conditions. Due to these requirements, we define two types of objective both based on deviations from target landing times.

The ALP is a highly dynamic problem, that is new aircraft are continuously appearing in the radar range while some are landing. However this research focuses on the static (off-line) case of the problem. This means that the number of aircraft that must be landed is fixed and all the aircraft characteristics are available at anytime.

2 Formulation

The mixed-integer zero-one formulation of the static multiple runway ALP presented in this work is based on the formulations presented in Beasley et al. (2000, 2001, 2004). A significant feature of this formulation is that separation times that must elapse between aircraft landings are not aircraft type dependent but they are uniquely defined for each pair of aircraft. For instance, an Upper-Medium aircraft landing after a Heavy aircraft can require a different separation time from another Upper-Medium aircraft landing after another Heavy aircraft. Practically, it enables other possible time requirements depending on flight characteristics and approach corridors used to be taken into account in addition to the separation time solely based on the aircraft type.

This formulation contains three main sets of constraints. The time window constraints as introduced earlier make sure that each aircraft land within its time window, that is between its earliest landing time and its latest landing time. The separation time constraints, also introduced earlier, make sure that for a pair of aircraft the appropriate separation time given the landing order of the aircraft and their allocated runway is satisfied. The last set is runway allocation constraints, they ensure that exactly one runway is assigned to each aircraft waiting

to land and that the aircraft's runway assignment is consistent when aircraft must land on the same runway.
Two different objectives are considered along with these constraints. The first objective is represented with a linear (or linearisable) function and is to make all aircraft land as close as possible to their target landing times. This objective could be applied when a reasonable traffic density holds. The second objective is represented with a non-linear function and is to make all aircraft land as early as possible. This objective is more likely to be applied when a high traffic density occurs and efficient and effective use of the runway capacity must be made.
There are a number of extensions to the formulation of the ALP presented that are worthy of note. All of these extensions relate to introducing runway dependence into the problem. In the formulation the earliest, target and latest times were assumed to be independent of the runway chosen for landing. In practice this may not be so. In particular aircraft close to the airport may well have significantly different earliest times for different runways. We present how to easily extend our formulation to incorporate runway dependence for these times. Similarly the separation time that must elapse between two aircraft landing on the same runway is assumed to be independent of the runway assigned. In practice this may not be so. For example noise abatement procedures may require a different separation time between aircraft if they land on one runway instead of another runway. We also present how to easily extend our formulation to incorporate runway dependence for separation times. A last extension that is considered is related to runway restrictions. It can happen that a given aircraft cannot make use of certain runways. For instance a runway may be too short for a Heavy-type aircraft to land on. This is easily dealt with in our formulation by setting the appropriate decision variable to the required value.
As mentioned previously the formulation presented in this work and all its extensions are developed for the static case of the multiple runway ALP. The set of aircraft waiting to land is known beforehand and no new incoming aircraft are considered.

3 Population Heuristics

The problem of scheduling aircraft landings at airports has been studied using various approaches (e.g. simulation, dynamic programming, queuing theory, heuristics). However few authors (Abela et al., 1993, Beasley et al., 2001, Cheng et al., 1999, Ciesielski and Scerri, 1997, 1998, Ernst et al., 1999, Hansen, 2004) make use of a population-based heuristic approach to solve this problem.
Genetic Algorithms are the most widely known Population Heuristic (PH). Such heuristics are based on the principles of selection and mutation, the main concepts of Darwin's theory of evolution. Given a population of individuals, those best adapted to the environment are said to be fitter and have a better chance to survive. If these fitter individuals mate together, mixing their genetic material, and the offspring created enter the population eliminating less fit individuals, then the population will evolve to become better adapted to its environment.
PHs mirror evolution in performing manipulations on individuals which represent possible solutions to the considered problem. Each individual is encoded using a set of chromosomes that define the problem's variables. The fitness of an individual is evaluated with respect to the quality of the solution it represents. An initial population of individuals is generated and operators that model genetic selection, mating and other processes are defined and applied to

the population individuals. The selection operator designates the individuals that will play the role of parents. The combination operator handles the generation of offspring. There can be a mutation operator that performs a slight perturbation to new individuals. Finally, some of the new individuals are chosen to be inserted into the current population in accordance with the population replacement scheme. This scheme also designates which of the current individuals must be removed from the population. A general framework for PHs —and therefore for GAs— is given in Beasley (2002) and presented below:

generate an initial population
repeat
    select individuals from the population to be parents
    create new individuals as combinations of selected parents
    optionally mutate the children
    select the children to insert into the population
    select the individuals to remove from the population
until termination, whereupon report the best solution encountered

In this research we develop two particular PHs: Scatter Search (SS) and the Bionomic Algorithm (BA) that have never been applied to the ALP before. They are used to generate new landing sequences and enable more effective solutions for the ALP to be found.
The difference between standard GAs and other PHs such as SS and BA used in this paper, relates to the strategies used to implement the steps in the above framework. Operators involved in GAs are mainly random procedures that could be applied to any type of problem. Other PHs involve deterministic procedures that can include problem specific knowledge. As there exist a large number of possible approaches to the steps involved in a PH, a major consideration is to choose appropriate approaches that will achieve good algorithmic performance. In particular, it is important to achieve a good balance between procedures that intensify the solution search and procedures that diversify this search.
Many of the elements involved in our implementation of SS and BA for the ALP are common to both algorithms, because either they correspond to basic features of PHs or they are directly related to the ALP. The major element that differs between SS and BA relates to the selection of parents. The elements of our heuristics are:

- representation of an individual
- evaluation of an individual
- generation of the initial population
- selection of parents
- generation of children
- duplication test
- local improvement of an individual
- population replacement

4 Computational Results

Computational results are presented for 13 problem instances involving from 10 aircraft to 500 aircraft and publicly available from OR-Library (see Beasley, 1990). The SS and BA presented in the previous section are implemented in C++ on a 2 GHz Pentium PC with 512 MB of memory. Varying the number of runways up to five and the objective adopted for these 13 instances means we consider, in total, 101 test problems. Results obtained with the FCFS solution are reported along with the best result obtained with both algorithms after ten executions for each test problem.

The very first result to mention is that the final solution provided by our algorithms is feasible in all of the 101 test problems considered.

For the non-linear objective, results quality is globally very high. The BA produces a better solution than SS for a larger number of test problems. Execution times for both algorithms increase as the number of aircraft increase (as we would expect). However for a given number of aircraft, there does not seem to be a marked increase in computation time as the number of runways increases. For the linear objective, computation times are larger than those for the NLO. The primary reason for this lies in the child improvement scheme adopted. For this objective, SS performs better than BA on a larger number of instances.

In order to provide insight into the computational effectiveness of our SS and BA heuristics, we implemented a basic local search heuristic (BLSH) for the static single runway ALP. This heuristic examines repositioning single aircraft and swapping the positions of pairs of aircraft in the landing sequence, in order to try and improve the solution. This movement of single aircraft in the landing sequence, and the swapping of the positions of pairs of aircraft, is repeated until no improvement in the landing sequence can be found. Results provided by BLSH are of high quality but at the expense of computation time.

Test problems relating to runway dependent landing times have been given in Cheng et al. (1999) and Hansen (2004). There are five test problems in these papers, all of which have runway dependent earliest times and one of which has runway restrictions. On average, both our SS and BA algorithms produce better quality solutions than the algorithms previously presented for this kind of test problems. Computationally the times for SS and BA are small for these five test problems.

5 Conclusions

In this paper we considered the static multiple runway aircraft landing problem. Two different objective functions were considered, a non-linear objective function and a linear objective function. We presented two population heuristic algorithms, scatter search and the bionomic algorithm, which have not been applied to the problem previously in the literature. Computational results were presented for a large number of problem instances involving up to 500 aircraft and 5 runways. These instances are much larger than those that have been considered by the majority of workers in the literature.

For the non-linear objective our results indicated that our bionomic algorithm out-performed our scatter search algorithm. For the linear objective the reverse was true with our scatter search algorithm out-performing our bionomic algorithm.

References

J. Abela, D. Abramson, M. Krishnamoorthy, A. De Silva and G. Mills, "Computing optimal schedules for landing aircraft", *Proceedings of the 12th National ASOR Conference* 71-90 (1993).

J.E. Beasley, "OR-Library:distributing test problems by electronic mail", *Journal of the Operational Research Society* 41, 1069-1072 (1990).

J.E. Beasley, "Population Heuristics," in *Handbook of applied optimization*, P.M. Pardalos and M.G.C. Resende (eds), 138-157, Oxford University Press, Oxford, 2002.

J.E. Beasley, M. Krishnamoorthy, Y.M. Sharaiha and D. Abramson, "Scheduling Aircraft Landings - The Static Case", *Transportation Science* 34, 180-197 (2000).

J.E. Beasley, M. Krishnamoorthy, Y.M. Sharaiha and D. Abramson, "Displacement problem and dynamically scheduling aircraft landings", *Journal of the Operational Research Society* 55, 54-64 (2004).

J.E. Beasley, J. Sonander and P. Havelock, "Scheduling aircraft landings at London Heathrow using a population heuristic", *Journal of the Operational Research Society* 52, 483-493 (2001).

V.H.L. Cheng, L.S. Crawford and P.K. Menon, "Air traffic control using genetic search techniques," presented at 1999 IEEE International Conference on Control Applications, Hawaii, August 1999.

V. Ciesielski and P. Scerri, "An anytime algorithm for scheduling of aircraft landing times using genetic algorithms", *Australian Journal of Intelligent Information Processing Systems* 4, 206-213 (1997).

V. Ciesielski and P. Scerri, "Real time genetic scheduling of aircraft landing times", *Proceedings of the 1998 IEEE international Conference on Evolutionary Computation* 360-364 (1998).

A.T. Ernst, M. Krishnamoorthy and R.H. Storer, "Heuristic and exact algorithms for scheduling aircraft landings", *Networks* 34, 229-241 (1999).

J.V. Hansen, "Genetic search methods in air traffic control", *Computers and Operations Research* 31, 445-459 (2004).